Iterated Walsh–Hadamard systems of the form

$x_k = (HD)^k x_0$

where H is a Walsh–Hadamard transform and D is a diagonal modulation matrix, define a useful class of structured high-dimensional dynamical systems. Although composed only of orthogonal mixing and sign modulation, repeated iteration generates a rich trajectory of correlated projections of the original input vector.

The system is computationally attractive because Walsh–Hadamard transforms require only additions and subtractions and can be implemented in O(nlogn) time. Despite this simplicity, repeated iteration rapidly transforms highly quantized inputs such as $\{-1,+1\}^n$ vectors into dense real-valued states with approximately Gaussian-like coordinate distributions.

Several algorithmic and application directions appear worth further exploration.

One promising area is reservoir and recurrent computing. Iterated HD systems naturally generate evolving internal state trajectories while preserving global signal energy. This may provide a lightweight alternative to large learned recurrent matrices in recurrent neural networks or echo-state systems:

$x_{t+1} = HDx_t + Bu_t$

The orthogonal structure helps maintain long-term information without the instability or collapse common in conventional recurrent architectures.

Another direction is structured random feature generation. The sequence

$\phi(x) = [x, HDx, (HD)^2x, ..., (HD)^k x]$

provides many related projections of the same input at very low cost. This may be useful for kernel approximation, approximate nearest-neighbor search, compressed embeddings, or low-parameter feature expansion systems. Unlike fully random projections, the transforms are reusable, deterministic, and hardware efficient.

The system may also have applications in quantized or binary neural computation. Binary input states are progressively transformed into dense floating-point-like trajectories through repeated orthogonal mixing. This effectively converts discrete symbolic structure into smooth continuous representations while preserving information globally. Such behavior may help bridge symbolic and statistical computation or improve low-bit neural architectures.

Another potentially important direction is trajectory-based computation. Most machine learning systems treat intermediate states as disposable layer activations. In iterated HD systems, however, the trajectory itself may contain useful information:

- orbit statistics,

- temporal correlations,
- recurrence structure,
- or projection agreement across iterations.

This suggests possible algorithms based on dynamical signatures rather than single static embeddings.

The approach also has strong connections to compressed sensing, hyperdimensional computing, and associative memory systems. Iterated orthogonal transforms distribute information broadly across coordinates while preserving norm, potentially improving robustness, redundancy, and retrieval behavior in high-dimensional representations.

An especially interesting research direction concerns modulation schedules. If the diagonal modulation varies over time,

$x_{k+1} = HD_k x_k$

then the resulting trajectories may exhibit quasi-random, ergodic, or low-discrepancy behavior. Carefully chosen modulation sequences based on parity functions, Weyl sequences, or Thue–Morse-like constructions may improve state-space coverage and avoid undesirable alignment with Walsh eigenspaces.

Finally, the combination of:

- reversibility,
- strong mixing,
- pseudo-randomization,
- and computational efficiency

is relatively uncommon in existing machine learning and signal-processing systems. This suggests that iterated HD dynamics may eventually prove useful not only as an optimization of existing architectures, but also as the basis for new classes of algorithms built around reversible high-dimensional projection flows.